



(12) **United States Patent**
Benhase et al.

(10) **Patent No.:** **US 9,262,088 B2**
(45) **Date of Patent:** **Feb. 16, 2016**

(54) **USE OF FLASH CACHE TO IMPROVE
TIERED MIGRATION PERFORMANCE**

12/0811 (2013.01); **G06F 12/0866** (2013.01);
G06F 12/0897 (2013.01); **G06F 2212/608**
(2013.01)

(71) Applicant: **INTERNATIONAL BUSINESS
MACHINES CORPORATION,**
Armonk, NY (US)

(58) **Field of Classification Search**
CPC **G06F 12/0811**; **G06F 3/0647**; **G06F 12/0897**
USPC **711/119**
See application file for complete search history.

(72) Inventors: **Michael T. Benhase**, Tucson, AZ (US);
Lokesh M. Gupta, Tucson, AZ (US)

(56) **References Cited**

(73) Assignee: **INTERNATIONAL BUSINESS
MACHINES CORPORATION,**
Armonk, NY (US)

U.S. PATENT DOCUMENTS

8,615,678 B1 * 12/2013 Madnani et al. 714/5.1
2010/0235582 A1 * 9/2010 Factor et al. 711/133
2011/0225117 A1 9/2011 Nakanishi et al.
2012/0059994 A1 3/2012 Montgomery et al.

OTHER PUBLICATIONS

(*) Notice: Subject to any disclaimer, the term of this
patent is extended or adjusted under 35
U.S.C. 154(b) by 228 days.

Choi, et al., Study of the performance impact of a cache buffer in
solid-state disks, Microprocessors and Microsystems, May 2011,
vol. 35 , No. 3 , pp. 359-369.

(21) Appl. No.: **14/066,147**

* cited by examiner

(22) Filed: **Oct. 29, 2013**

Primary Examiner — Jae Yu

(65) **Prior Publication Data**

(74) *Attorney, Agent, or Firm* — Griffiths & Seaton PLLC

US 2014/0208032 A1 Jul. 24, 2014

(57) **ABSTRACT**

Related U.S. Application Data

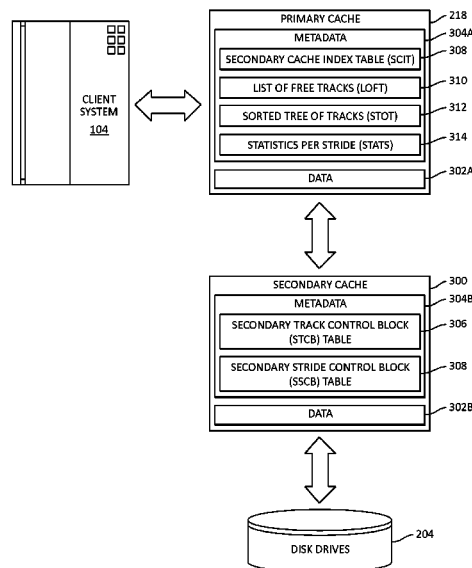
(63) Continuation of application No. 13/746,973, filed on
Jan. 22, 2013, now Pat. No. 9,032,153.

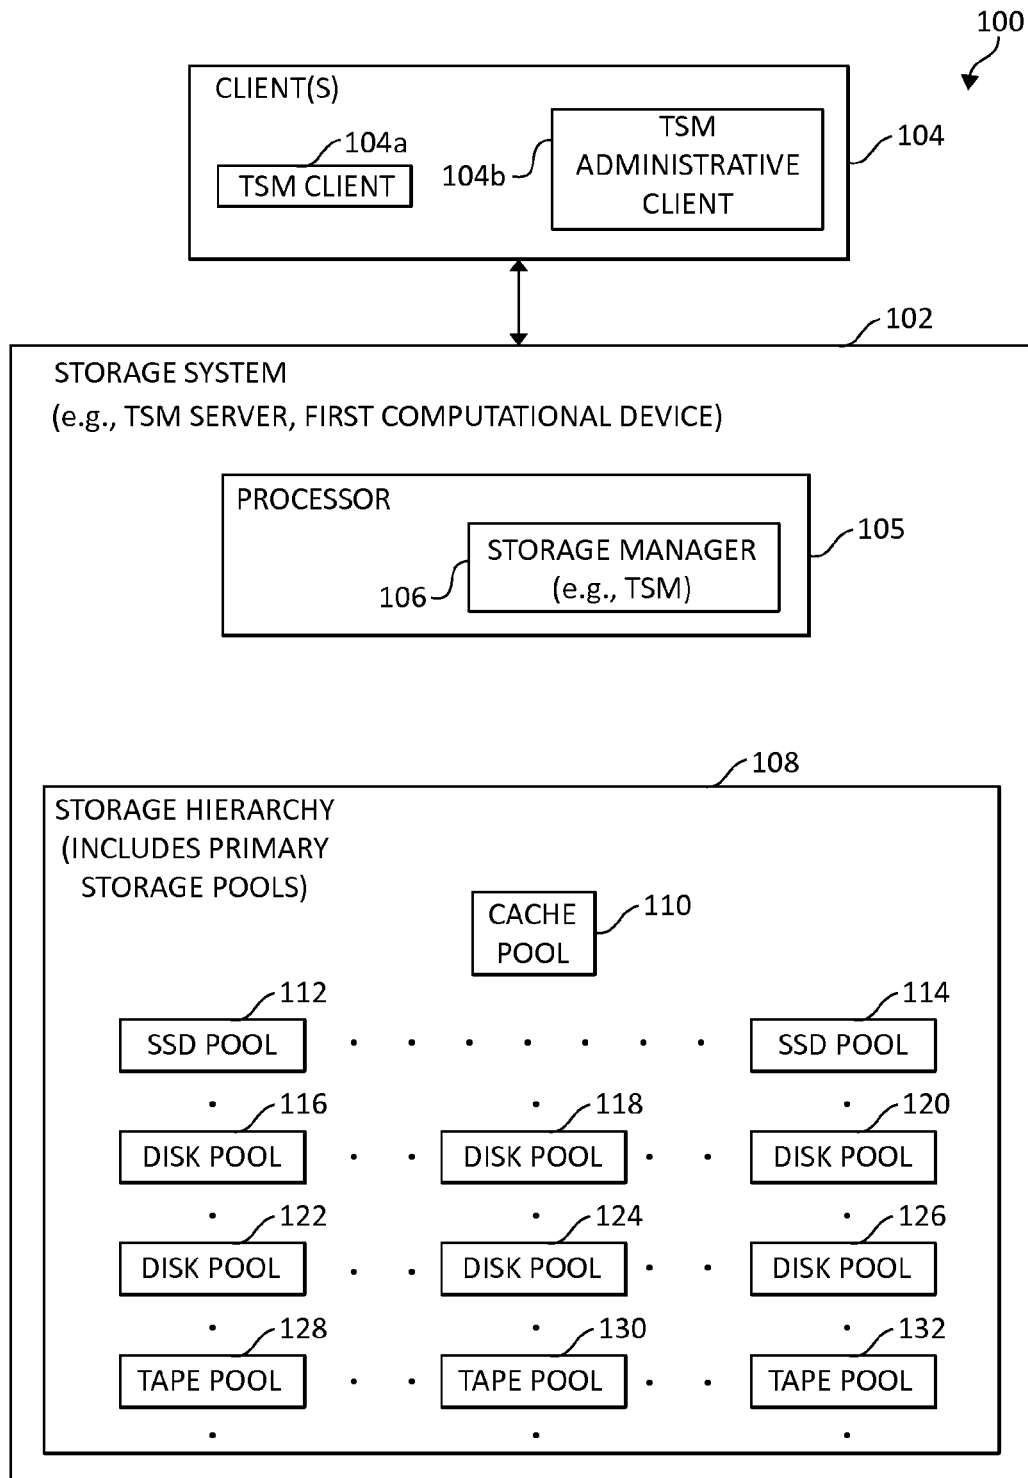
For data processing in a computing storage environment by a
processor device, the computing storage environment incor-
porating at least high-speed and lower-speed caches, and
tiered levels of storage, and at a time in which at least one data
segment is to be migrated from one level to another level of
the tiered levels of storage, a data migration mechanism is
initiated by copying data resident in the lower-speed cache
corresponding to the at least one data segment to be migrated
to a target on the another level, and reading remaining data,
not previously copied from the lower-speed cache, from a
source on the one level, and writing the remaining data to the
target.

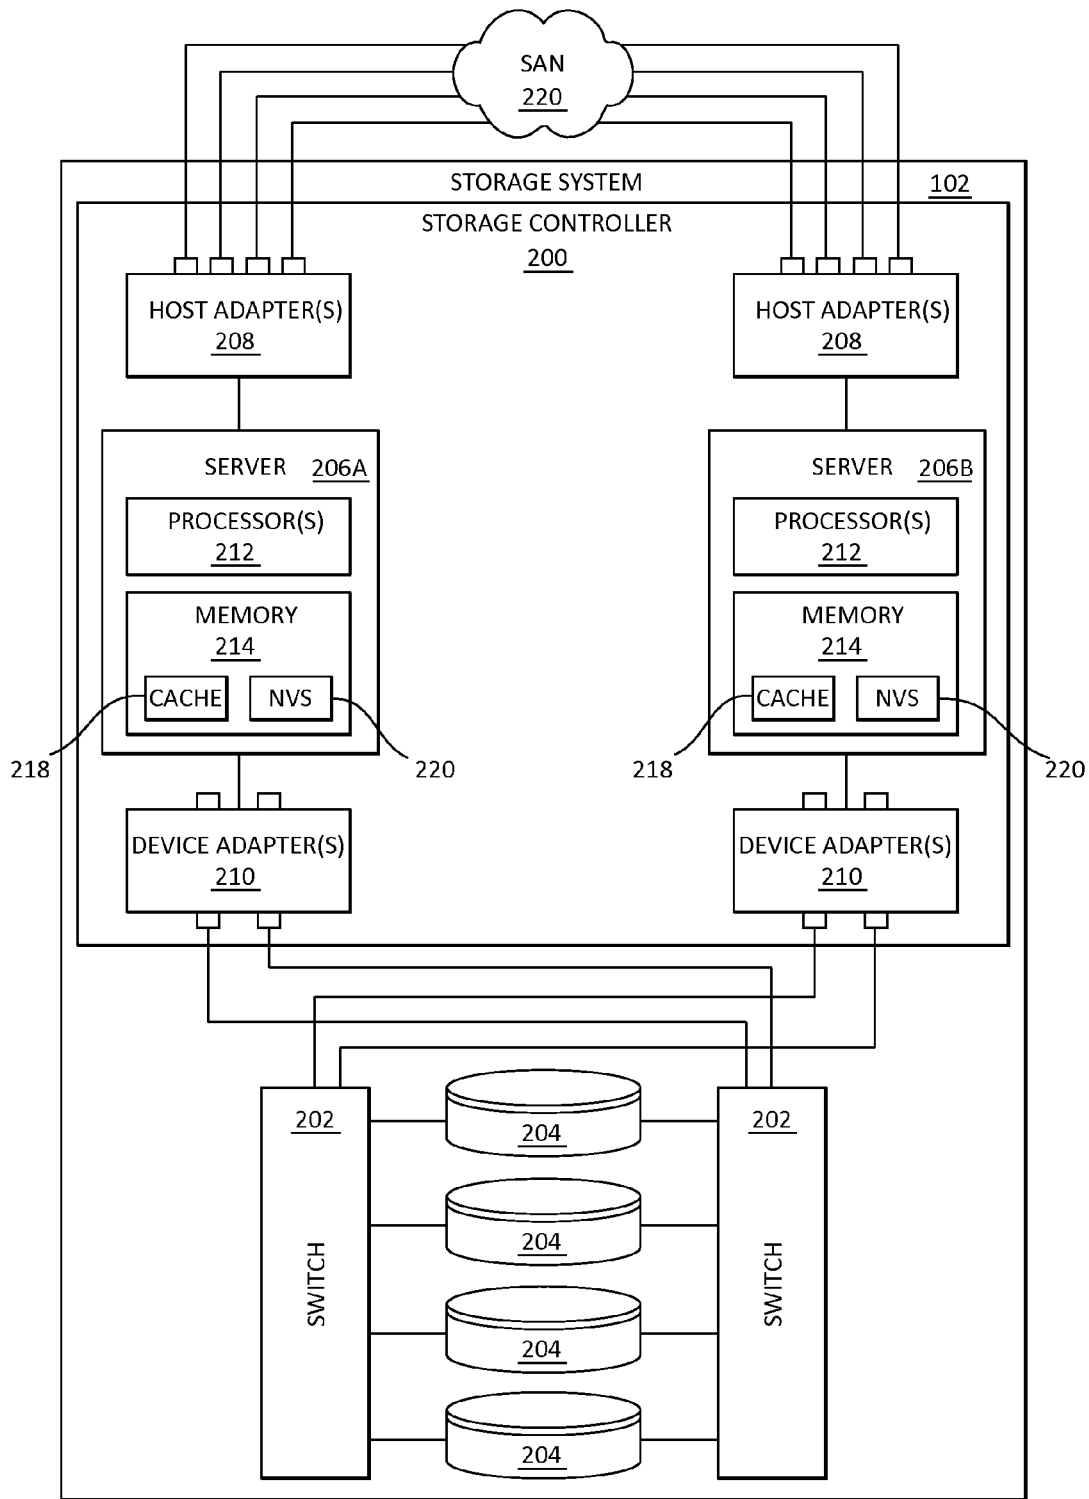
6 Claims, 5 Drawing Sheets

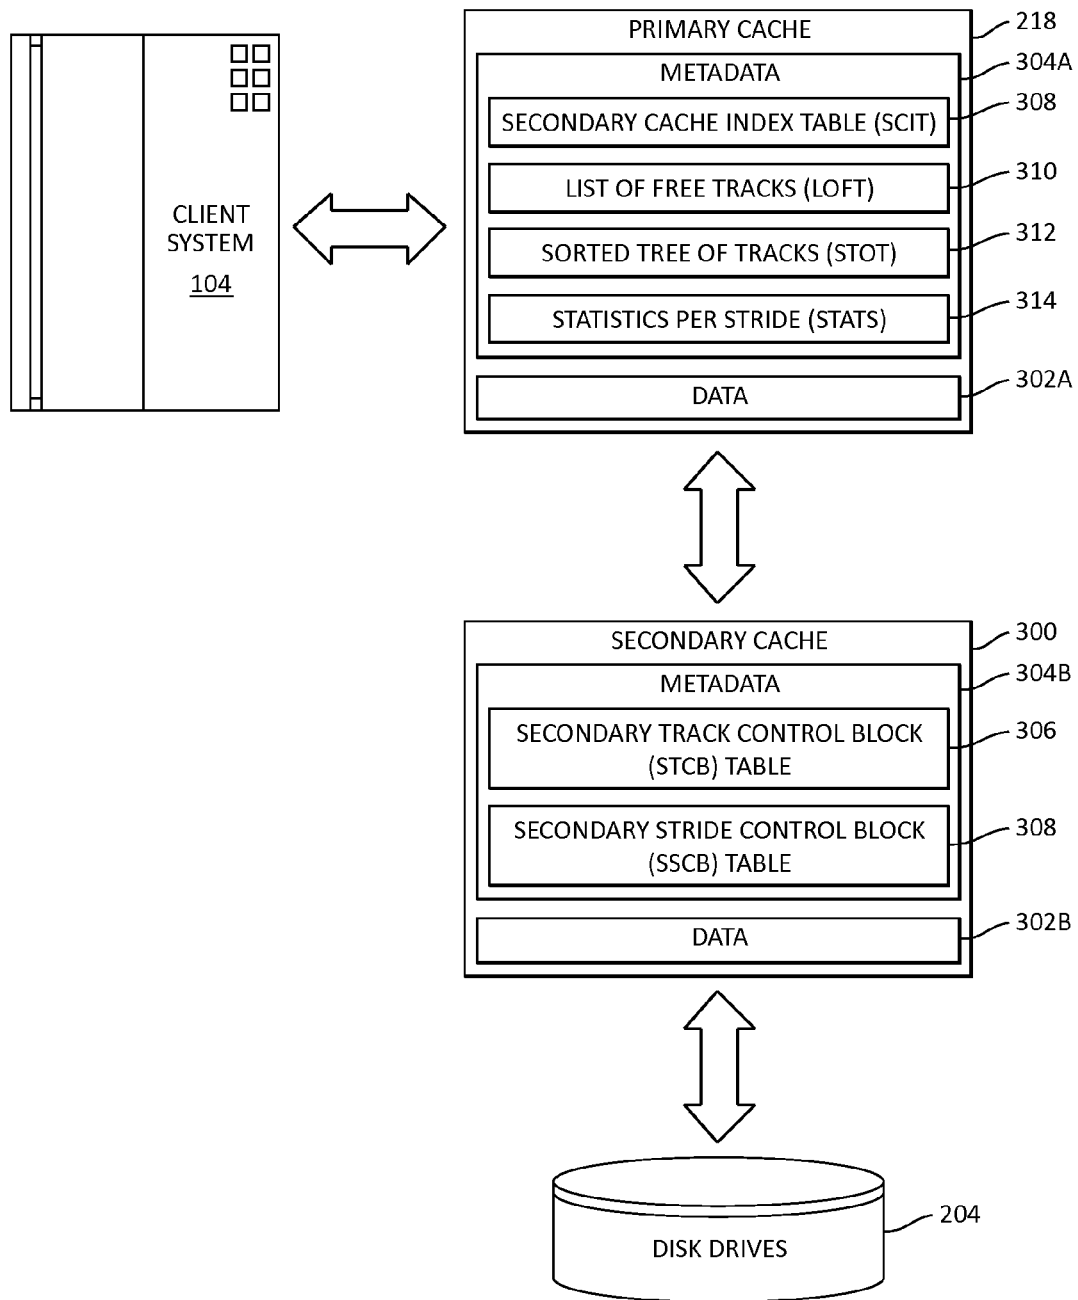
(51) **Int. Cl.**
G06F 12/10 (2006.01)
G06F 3/06 (2006.01)
G06F 12/08 (2006.01)

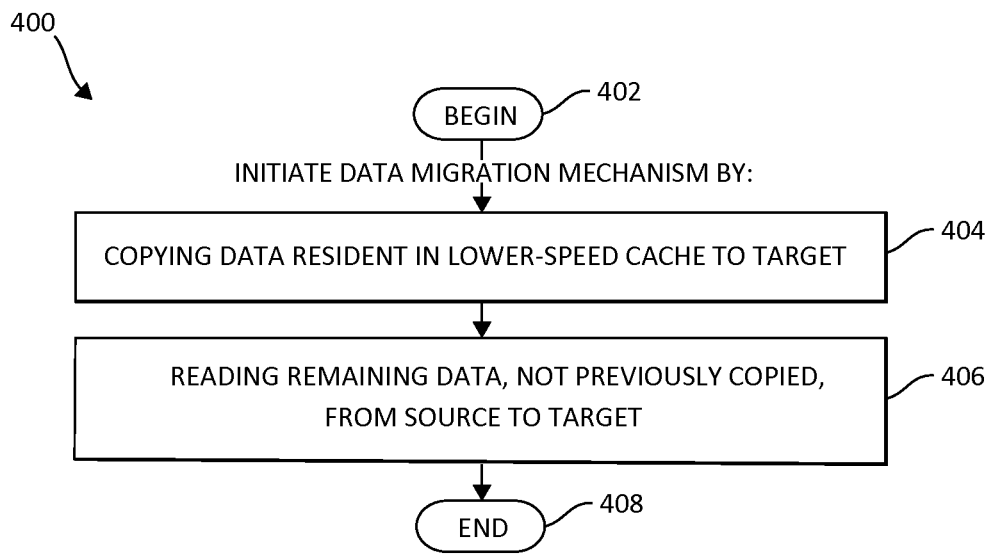
(52) **U.S. Cl.**
CPC **G06F 3/0647** (2013.01); **G06F 3/0608**
(2013.01); **G06F 3/0673** (2013.01); **G06F**

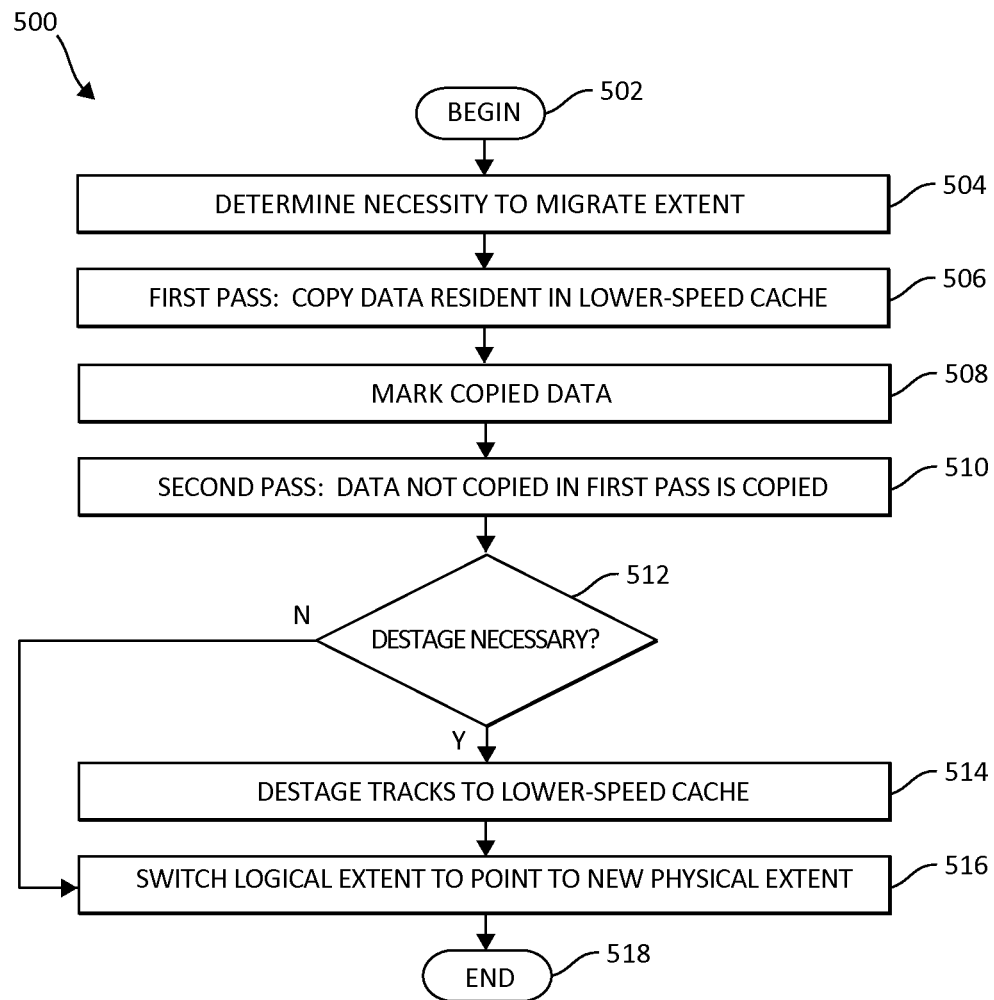


**FIG. 1**

**FIG. 2**

**FIG. 3**

**FIG. 4**

**FIG. 5**

1

USE OF FLASH CACHE TO IMPROVE TIERED MIGRATION PERFORMANCE

CROSS REFERENCE TO RELATED APPLICATIONS

This application is a Continuation of U.S. patent application Ser. No. 13/746,973, filed on Jan. 22, 2013.

FIELD OF THE INVENTION

The present invention relates generally to computers, and more particularly, to mechanisms for migrating data segments in a computing storage environment.

DESCRIPTION OF THE RELATED ART

In today's society, computer systems are commonplace. Computer systems may be in the field of computing, a "cache" typically refers to a small, fast memory or storage device used to store data or instructions that were accessed recently, are accessed frequently, or are likely to be accessed in the future. Reading from or writing to a cache is typically cheaper (in terms of access time and/or resource utilization) than accessing other memory or storage devices. Once data is stored in cache, it can be accessed in cache instead of re-fetching and/or re-computing the data, saving both time and resources.

SUMMARY OF THE DESCRIBED EMBODIMENTS

Caches may be provided as multi-level caches. For example, a caching system may include both a "primary" and "secondary" cache. When reading data, a computing system or device may first look for data in the primary cache and, if the data is absent, look for the data in the secondary cache. If the data is not in either cache, the computing system or device may retrieve the data from disk drives or other storage devices. When writing data, a computing system or device may write data to the primary cache. This data may eventually be destaged to the secondary cache or a storage device to make room in the primary cache.

Storage environments in computer systems may include so-called tiered storage architectures, which may include a variety of storage mediums such as enterprise hard disk drives (HDD), serial advanced technology attachment (SATA) disk drives, solid state drives (SSD), tape drives, and other devices. Data may be migrated between such devices. Placement of certain kinds of data in the appropriate medium may greatly enhance overall performance of the storage environment.

In an existing tiered migration mechanism, data is read from a source location and destaged to a target. During this step, data can still be destaged to the source, which if occurs is recorded in a bitmap. In a subsequent step, any new destages are withheld to the source while tracks are read that were previously destaged to the source and destaged to the target. This mechanism results in additional time being expended while new destages are withheld, and while the previously destaged tracks are re-read and destaged from the source to the target. A need exists for a mechanism that effectively reduces this additional time, among other factors, to increase storage performance.

Accordingly, and in view of the foregoing, various exemplary method, system, and computer program product

2

environment incorporating at least high-speed and lower-speed caches, and tiered levels of storage, are provided. At a time in which at least one data segment is to be migrated from one level to another level of the tiered levels of storage, a migration mechanism is initiated by copying data resident in the lower-speed cache corresponding to the at least one data segment to be migrated to a target on the another level, and reading remaining data, not previously copied from the lower-speed cache, from a source on the one level, and writing the remaining data to the target. In addition to the foregoing exemplary embodiment, various other system and computer program product embodiments are provided and supply related advantages.

The foregoing summary has been provided to introduce a selection of concepts in a simplified form that are further described below in the Detailed Description. This Summary is not intended to identify key features or essential features of the claimed subject matter, nor is it intended to be used as an aid in determining the scope of the claimed subject matter. The claimed subject matter is not limited to implementations that solve any or all disadvantages noted in the background.

BRIEF DESCRIPTION OF THE DRAWINGS

In order that the advantages of the invention will be readily understood, a more particular description of the invention briefly described above will be rendered by reference to specific embodiments that are illustrated in the appended drawings. Understanding that these drawings depict embodiments of the invention and are not therefore to be considered to be limiting of its scope, the invention will be described and explained with additional specificity and detail through the use of the accompanying drawings, in which:

FIG. 1 is a high-level block diagram showing an exemplary tiered storage environment in which mechanisms in accordance with the present invention may be implemented;

FIG. 2 is a high-level block diagram showing an exemplary storage system in which mechanisms in accordance with the present invention may be implemented;

FIG. 3 is a high-level block diagram showing one embodiment of a multi-level cache in accordance with the invention;

FIG. 4 is a flow chart diagram showing an exemplary method for using flash cache to improve storage migration performance, in accordance with one embodiment of the present invention; and

FIG. 5 is a flow chart diagram showing an additional exemplary method for using flash cache to improve storage migration performance, again in accordance with one embodiment of the present invention.

DETAILED DESCRIPTION OF CERTAIN EMBODIMENTS

As previously mentioned, in an existing tiered migration mechanism, data is read from a source location and destaged to a target. During this step, data can still be destaged to the source, which if occurs is recorded in a bitmap. In a subsequent step, any new destages are withheld to the source while tracks are read that were previously destaged to the source and destaged to the target. This mechanism results in additional time being expended while new destages are withheld. Any new writes from the host to these data segments will need to wait for the destages to finish, which may take a significant amount of time, since it may take such time to copy data from a source to target location (e.g., while the previously destaged tracks are re-read and destaged from the source to the target). In addition, this current mechanism does not examine data in

cache when migration operations are performed. Such mechanism starts at a beginning of a data segment (e.g., an extent) and proceeds to the end in order. If tracks are being maintained in cache, alternatively, these tracks may be copied first, before the tracks get aged out of cache.

In view of this scenario, the mechanisms of the present invention utilize primary and secondary cache architectures to allow for more direct involvement of such caches in the data migration process. As a result, for example, destages (i.e., resources in the computing storage environment) are not held while other storage operations (such as the aforementioned reads and writes of the data being migrated) occur. For example, destage conflicts which may occur during data segment migration operations are in progress may be significantly reduced. Accordingly, overall storage performance and efficiencies are enhanced as will be further described.

For the purposes of this disclosure, the phrase “secondary cache” is used to refer to any cache (including, for example, L2 or L3 cache) that resides between a primary cache and a storage device, such as a disk drive, tape drive, or the like.

Turning now to FIG. 1, a block diagram of a tiered computing storage environment 100 in accordance with certain embodiments is illustrated. The computing storage environment 100 includes a first computational device, such as, a storage system 102, coupled to one or more computational devices, such as, clients 104. In certain embodiments, the storage system 102 and the clients 104 may comprise any suitable computational device, including those presently known in the art, such as, a personal computer, a workstation, a mainframe, a midrange computer, a network appliance, a palm top computer, a telephony device, a blade computer, a hand held computer, etc.

In some embodiments, a storage manager 106, such as, the Tivoli® Storage Manager® (TSM) product marketed by International Business Machines (IBM®) Corporation, may be used for securely storing and managing data segments according to aspects of the present invention. The storage manager 106 may execute in a storage management server, such as a TSM server 102 or elsewhere. In one embodiment, the storage manager is operable by and/or in conjunction with processor device 105 as shown. One of ordinary skill in the art will appreciate that various other configurations of the processor 105, storage manager 106, and related additional processing and/or memory components are contemplated. IBM, Tivoli, and Tivoli Storage Manager are trademarks or registered trademarks of IBM Corporation.

The TSM may provide data storage services to clients, such as TSM clients 104a, for management of data. The TSM server 102 may store files sent to the TSM server 102 by one or more TSM clients 104a, 104b. The storage manager 106 and/or processor device 105 may allow a system administrator to configure storage pools, where a storage pool comprises a group of devices used for storing data received from the TSM clients 104a, 104b. Storage pools are used as targets for store operations from the TSM clients 104a, 104b, and are referenced in TSM server policies and other constructs for processing.

As shown, a variety of storage devices may be organized into a storage hierarchy. Storage media within the storage hierarchy may thus be grouped into data structures referred to herein as storage pools. The storage hierarchy may be organized to correspond with one or more metrics, such as a performance metric including write or read speeds. The storage hierarchy 108 as shown may be organized such that the top of the hierarchy may include a cache pool 110 having a highest amount or quality of a particular performance metric. Below the cache pool 110, a number of solid state drive (SSD)

class devices may be organized into SSD pools by the same, similar, or other metrics (e.g., SSD pools 112 and 114).

Below the SSD pools 112 and 114, a first tier of disk pools (e.g., disk pools 116, 118, and 120) may be then organized. As one of ordinary skill in the art will appreciate, disk pools 116, 118, and 120 may include a variety of disk devices such as pools of enterprise disk drives, SATA disk drives, disk devices configured in a particular redundant array of independent disks (RAID) configuration, and the like.

The first tier of disk pools may be located above a second tier of disk pools (e.g., pools 122, 124, and 126) by virtue of exhibiting a greater amount, stronger attribute or attributes, or quality of the performance metric. Below the second tier of disk pools, an additional tier of tape pools (e.g., tape pools 128, 130, and 132) may then be organized. Various considerations for the organization of such storage hierarchies 108 may be apparent to one of ordinary skill in the art. In one embodiment, the system administrator may assist in performing such configurations in the storage hierarchy 108 by inputs to the TSM administrative client 104b or another mechanism. While tape pools 128, 130, and 132 are shown within the storage hierarchy 108 as shown, it should be noted that generally such tape pools are found in a storage subsystem external to those pools ranking higher in the hierarchy.

Referring to FIG. 2, one embodiment of a storage system 102 containing an array of hard-disk drives 204 and/or solid-state drives 204 is illustrated. The internal components of the storage system 102 are shown since the caching system may, in certain embodiments, be implemented within such a storage system 102, although the caching system may also be applicable to other storage systems 102. As shown, the storage system 102 includes a storage controller 200, one or more switches 202, and one or more storage devices 204 such as hard disk drives 204 or solid-state drives 204 (such flash-memory-based drives 204). The storage controller 200 may enable one or more clients 104 (e.g., open system and/or mainframe servers 104) to access data in the one or more storage devices 204. The clients 104 (e.g., FIG. 1) may be accessible through a Storage Area Network (SAN) 220 as shown.

In selected embodiments, the storage controller 200 includes one or more servers 206. The storage controller 200 may also include host adapters 208 and device adapters 210 to connect the storage controller 200 to host devices 106 and storage devices 203, 204, respectively. Multiple servers 206a, 206b may provide redundancy to ensure that data is always available to connected hosts 106. Thus, when one server 206a fails, the other server 206b may remain functional to ensure that I/O is able to continue between the clients 104 and the storage devices 204. This process may be referred to as a “failover.”

One example of a storage system 102 having an architecture similar to that illustrated in FIG. 2 is the IBM® DS8000™ enterprise storage system. The DS8000™ is a high-performance, high-capacity storage controller providing disk storage that is designed to support continuous operations. The DS8000™ series models may use IBM's POWER5™ servers 206a, 206b which may be integrated with IBM's virtualization engine technology. Nevertheless, the caching system disclosed herein is not limited to the IBM® DS8000™ enterprise storage system, but may be implemented in any comparable or analogous storage system 110, regardless of the manufacturer, product name, or components or component names associated with the system 110. Furthermore, any system that could benefit from one or more embodiments of the invention is deemed to fall within the

scope of the invention. Thus, the IBM® DS8000™ is presented only by way of example and is not intended to be limiting.

In selected embodiments, each server **206** may include one or more processors **212** (e.g., n-way symmetric multiprocessors) and memory **214**. The memory **214** may include volatile memory (e.g. RAM) as well as non-volatile memory (e.g., ROM, EPROM, EEPROM, hard disks, flash memory, etc.). The volatile memory and non-volatile memory may, in certain embodiments, store software modules that run on the processor(s) **212** and are used to access data in the storage devices **204**. The servers **206** may host at least one instance of these software modules. These software modules may manage all read and write requests to logical volumes in the storage devices **204**.

In selected embodiments, the memory **214** may include a cache **218**. Whenever a client **104** (e.g., an open system or mainframe server **104**) performs a read operation, the server **206** that performs the read may fetch data from the storage devices **204** and save it in its cache **218** in the event it is required again. If the data is requested again by a client **104**, the server **206** may fetch the data from the cache **218** instead of fetching it from the storage devices **204**, saving both time and resources. Similarly, when a client **104** performs a write, the server **206** that receives the write request may store the write in its cache **218**, and destage the write to the storage devices **204** at a later time. When a write is stored in cache **218**, the write may also be stored in non-volatile storage (NVS) **220** of the opposite server **206** so that the write can be recovered by the opposite server **206** in the event the first server **206** fails.

Referring to FIG. 3, while continuing to refer generally to FIG. 2, as previously mentioned, a storage system **102** may include both hard disk drives **204** and solid-state drives (SSDs) **204**, such as flash-memory-based drives **204**. The I/O performance of SSDs **204** or other types of solid-state memory is typically far higher than the I/O performance of hard disk drives **204**. Because of the higher I/O performance, the solid-state drives **204** may, in certain embodiments, be used to provide a large secondary cache **300** between the primary cache **218** and the hard disk drives **204**. This large secondary cache **300** may significantly improve the I/O performance of the storage system **102**, and may be referred to herein as “Flash Cache.” Herein, the primary cache may be referred to as a high-speed or higher-speed cache (as it typically has access to the fastest Dynamic Read Only Memory or DRAM architectures), and the secondary, Flash Cache may be referred to as a low-speed or lower-speed cache (in comparison to the primary, DRAM cache), although this terminology is not intended to be limiting in any way.

Using the secondary cache **300**, if a read request is received by a server **206** (e.g., FIG. 2), the server **206** may initially look for data in the primary cache **218** and, if the data is not present, look for the data in the secondary cache **300** (residing in the solid-state drives **204**). If the data is not available in either cache, the server **206** may retrieve the data from the disk drives **204**. Similarly, when writing data, a server **206** may initially write the modified data to the primary cache **218**. This modified data may eventually be destaged to the secondary cache **300** to make room in the primary cache **218**. This data may then be destaged to the disk drives **204** to make space in the secondary cache **300**, as needed.

In certain embodiments, the secondary cache **300** may be sized to provide about one to twenty percent, or in other embodiments about five percent of the total storage capacity of the storage system **102**. Thus, for a storage system **102** that contains about ten terabytes (TB) of storage (from both hard

disk drives **204** and solid state drives **204**), about 0.5 TB of this storage space may be used as a secondary, “Flash” cache **300**. Such a large amount of secondary cache **300** may allow data to be destaged from the secondary cache **300** far less frequently than conventional primary or secondary caches. As an example, a very large secondary cache **300** could store writes for an entire day without having to destage the writes to the disk drives **204**. The writes could then be destaged at night or during a period of relative inactivity. Cache management algorithms may be redesigned to efficiently utilize the additional space in the secondary cache **300** and take advantage of the performance improvements that are possible using a large secondary cache **300**.

As shown in FIG. 3, each cache **218**, **300** may store data **302a**, **302b** and metadata **304a**, **304b**. The data **302a**, **302b** may be stored in the form of tracks. Each track in the secondary cache **300** may have a secondary track control block (STCB) associated therewith. The STCB may also be referred to herein as Cache Flash Control Block (CFCB). Along with other information, the STCB for each track may include a pointer to the next track in the chain, information indicating whether the track is free or in-use, as well as information indicating which sectors in the track have, been modified. In certain embodiments, the STCBs for all the tracks may be stored in an STCB table **306** stored in the secondary cache **300** as shown, or elsewhere.

In addition, each track in the secondary cache **300** may have a secondary stride control block (SSCB) associated therewith. The SSCB, like the STCB may include diagnostic and/or statistical information, but instead relating to strides (groups of tracks) stored in the secondary cache **300**. The SSCB may also be referred to herein as Cache Flash Element (CFE). In certain embodiments, the SSCBs for all the strides may be stored in an SSCB table **308** stored in the secondary cache **300** as shown, or elsewhere.

Similarly, the primary cache **218** may also store metadata **304a** associated with the secondary cache **300**. For example, the primary cache **218** may store a secondary cache index table (SCIT) **308** that provides a directory for tracks in the secondary cache **300**. In certain embodiments, the SCIT **308** is essentially a hash table with a constant hash function. To locate a specific track in the SCIT **308**, the hash function may convert a track identifier (e.g., a track number) to a hash value. This hash value may then be looked up in the SCIT **308** to find the STCB for the track. Alternatively, the SCIT **308** could be incorporated into a cache directory of the primary cache **218**, thereby providing a single hash table that stores tracks for both the primary and secondary caches **218**, **300**. In selected embodiments, the SCIT **308** is kept exclusively in the primary cache **218**. The SCIT **308** may be built or rebuilt (in the event of a failover, fallback, or initial microcode load (IML)) by reading the STCB table **306** in the secondary cache **300**.

In certain embodiments, the primary cache **218** may also store a list of free tracks (LOFT) **310** that indicates which tracks in the secondary cache **300** are free (i.e., unoccupied). This list **310** may be used to locate free space in the secondary cache **300** in order to destage data from the primary cache **218** to the secondary cache **300**. In selected embodiments, inserting or removing tracks from the LOFT **310** may be performed in a log structured manner. For example, tracks may be inserted at the end of the LOFT **310** and deleted from the front of the LOFT **310**. The LOFT **310** may be kept exclusively in the primary cache **218** and may be built or rebuilt by reading the STCB table **306** in the secondary cache **300**.

The primary cache **218** may also store a sorted tree of tracks (STOT) **312** that sorts the tracks by “trackid” or some other indicator. The STOT **312** may be used to minimize seek

time (on the disk drives **204**) when destaging tracks from the secondary cache **300** to the disk drives **204**. The STOT **312** may be kept exclusively in the primary cache **218** and may be built or rebuilt by reading the STCB table **306** in the secondary cache **300**.

The primary cache **218** may also store statistics per stride (STATS) **314** for each stride having one or more tracks in the secondary cache **300**. A “stride” refers to a set of logically sequential data that might be segmented across multiple disks combined with additional parity information as is for example used in a RAID-5 (redundant array of inexpensive disks) configuration. In general, the STATS **314** may be used to determine which tracks require the least number of disk operations (“disk ops”) to destage from the secondary cache **300** to the disk drives **204**. In general, the destage penalty for a track will be less where more tracks are present in a stride. When selecting, tracks to destage, tracks requiring the least number of disk ops may be destaged first to minimize resource utilization. In selected embodiments, the STATS **314** may store information such as the number of tracks that are present in the secondary cache **300** for each stride, and the number of disk ops required to destage a track in a stride. In certain embodiments, the STATS **314** may store a “recency” bit for each stride. The recency bit may be incremented each time an eviction process passes through a stride. The recency bit may be reset each time a track is added to a stride. The recency bit may be used to keep strides in the secondary cache **300** that are actively being written to. The STATS **314** may be kept exclusively in the primary cache **218** and may be built or rebuilt by reading the STCB table **306** in the secondary cache **300**.

The metadata **304a**, **304b** described above may be structured and stored in various different ways and is not limited to the illustrated structure or organization. The metadata **304a**, **304b** is provided by way of example to show one technique for storing and structuring the metadata **304a**, **304b**. For example, in certain embodiments, the data and metadata may be stored together in the secondary cache **300** in a circular log-structured array. Other methods for structuring and storing metadata **304a**, **304b** may be used and are encompassed within the scope of the invention.

As previously mentioned, one advantage of a large secondary cache **300** is that data can be destaged from the secondary cache **300** far less frequently than conventional secondary caches. This may enable more data to accumulate in the secondary cache **300** before it is destaged to the disk drives **204**. Accordingly, in selected embodiments, an apparatus and method in accordance with the invention may be configured to wait for full strides of data to accumulate and coalesce in the secondary cache **300** before the data is destaged to the disk drives **204**. As explained above, this may minimize the number of disk ops required to destage data from the secondary cache **300** to the disk drives **204**, thereby improving overall system performance.

Referring to FIG. 4, a first exemplary method for migration of data segments using secondary cache (Flash Cache), in a computing environment having dual lower and higher speed levels of cache, and tiered levels of storage, is illustrated. In the illustrated embodiment, the secondary cache is represented as the lower speed level of cache, and the higher speed cache may be implemented in the storage controller as DRAM cache as in a previous exemplary illustration. Method **400** begins (step **402**) as a data movement mechanism is initiated. The data movement mechanism is adapted for, implementing Flash Cache (for example, in 1-128 blocks) to perform data copy/migration (for example in 10 MB and 1 GB segments). Flash Cache is implemented to perform data

migration/copy functionality since the Flash Cache is physically and logically prioritized higher than other tiered levels of storage and can accommodate all Input/Output (I/O).

To avoid destage conflicts that may result in a DRAM demote to Flash Cache if a data segment (e.g., stride) is being destaged in progress, an out of order migration operation (similar to a RAID rebuild) for Read Miss, and Write (Flash Cache) may be incorporated. Accordingly, in step **404**, data resident in the lower-speed (Flash Cache) is copied to the target, and in step **406**, the remaining data, not previously copied, is read from the source and copied to the target (step **408**). The method **400** then ends (step **408**).

The data migration mechanism implemented using Flash Cache may be further adapted for implementing a preference of data movement of the partial data segments to the lower speed cache (again, e.g., Flash Cache) based on several metrics. Two possible such metrics are the amount of holes and data “hotness,” or a data heat metric. One objective of the preference of data movement previously described is to free up more space in the higher speed cache with a subsequent destage operation (more data being destaged) and coalesce into a single write to the secondary, lower speed cache.

FIG. 5, following, illustrates an exemplary method **500** of operation of the data movement mechanism as previously implemented in FIG. 4. Method **500** begins (step **502**) with a determination of a need to migrate a data segment (e.g., extent) by the tiered storage (step **504**). In a first pass, the tiered storage passes through the lower-speed cache and copies data resident in the cache to the target extent (step **506**). The data that has been copied is marked (for example with a bit or recorded in a bitmap) (step **508**).

In a subsequent step, a second pass is completed. In this pass, the tiered storage system copies that data which has not been copied in the first pass (step **510**). The source file is read and then the target file is written to. In this pass, if it is determined that data must be destaged from DRAM cache (higher-speed) to make space for new writes (step **512**), then the data segments are destaged to Flash Cache (lower-speed) (step **514**). This significantly reduces the destage conflicts since destages are not held off for the entire pass.

Once the second pass is complete, the logical extent is switched to point to the new physical extent (or whichever data segment is undergoing migration) (step **516**), and the method **500** ends (step **518**).

In some cases, for example if the Flash Cache is currently full of modified data, or otherwise inaccessible, it may not be possible at a particular point in time to destage to Flash Cache. In this case, the tiered storage system may switch to an alternative method of destaging in which Flash Cache is not implemented (and destages are held, for example, until data is migrated).

As will be appreciated by one skilled in the art, aspects of the present invention may be embodied as a system, method or computer program product. Accordingly, aspects of the present invention may take the form of an entirely hardware embodiment, an entirely software embodiment (including firmware, resident software, micro-code, etc.) or an embodiment combining software and hardware aspects that may all generally be referred to herein as a “circuit,” “module” or “system.” Furthermore, aspects of the present invention may take the form of a computer program product embodied in one or more computer readable medium(s) having computer readable program code embodied thereon.

Any combination of one or more computer readable medium(s) may be utilized. The computer readable medium may be a computer readable signal medium or a computer readable storage medium. A computer readable storage medium

may be, for example, but not limited to, an electronic, magnetic, optical, electromagnetic, infrared, or semiconductor system, apparatus, or device, or any suitable combination of the foregoing. More specific examples (a non-exhaustive list) of the computer readable storage medium would include the following: an electrical connection having one or more wires, a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM or Flash memory), an optical fiber, a portable compact disc read-only memory (CD-ROM), an optical storage device, a magnetic storage device, or any suitable combination of the foregoing. In the context of this document, a computer readable storage medium may be any tangible medium that may contain, or store a program for use by or in connection with an instruction execution system, apparatus, or device.

Program code embodied on a computer readable medium may be transmitted using any appropriate medium, including but not limited to wireless, wired, optical fiber cable, RF, etc., or any suitable combination of the foregoing. Computer program code for carrying out operations for aspects of the present invention may be written in any combination of one or more programming languages, including an object oriented programming language such as Java, Smalltalk, C++ or the like and conventional procedural programming languages, such as the "C" programming language or similar programming languages. The program code may execute entirely on the user's computer, partly on the user's computer, as a stand-alone software package, partly on the user's computer and partly on a remote computer or entirely on the remote computer or server. In the latter scenario, the remote computer may be connected to the user's computer through any type of network, including a local area network (LAN) or a wide area network (WAN), or the connection may be made to an external computer (for example, through the Internet using an Internet Service Provider).

Aspects of the present invention are described above with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems) and computer program products according to embodiments of the invention. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, may be implemented by computer program instructions. These computer program instructions may be provided to a processor of a general purpose computer, special purpose computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor of the computer or other programmable data processing apparatus, create means for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks.

These computer program instructions may also be stored in a computer readable medium that may direct a computer, other programmable data processing apparatus, or other devices to function in a particular manner, such that the instructions stored in the computer readable medium produce an article of manufacture including instructions which implement the function/act specified in the flowchart and/or block diagram block or blocks. The computer program instructions may also be loaded onto a computer, other programmable data processing apparatus, or other devices to cause a series of operational steps to be performed on the computer, other programmable apparatus or other devices to produce a computer implemented process such that the instructions which execute on the computer or other programmable apparatus

provide processes for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks.

The flowchart and block diagram in the above figures illustrate the architecture, functionality, and operation of possible implementations of systems, methods and computer program products according to various embodiments of the present invention. In this regard, each block in the flowchart or block diagrams may represent a module, segment, or portion of code, which comprises one or more executable instructions for implementing the specified logical function(s). It should also be noted that, in some alternative implementations, the functions noted in the block might occur out of the order noted in the figures. For example, two blocks shown in succession may, in fact, be executed substantially concurrently, or the blocks may sometimes be executed in the reverse order, depending upon the functionality involved. It will also be noted that each block of the block diagrams and/or flowchart illustration, and combinations of blocks in the block diagrams and/or flowchart illustration, may be implemented by special purpose hardware-based systems that perform the specified functions or acts, or combinations of special purpose hardware and computer instructions.

While one or more embodiments of the present invention have been illustrated in detail, one of ordinary skill in the art will appreciate that modifications and adaptations to those embodiments may be made without departing from the scope of the present invention as set forth in the following claims.

What is claimed is:

1. A method for data processing in a computing storage environment, the computing storage environment incorporating at least high-speed and lower-speed caches, and tiered levels of storage, comprising:

at a time in which at least one data segment is to be migrated from one level to another level of the tiered levels of storage, initiating a data migration mechanism by:

- copying data resident in the lower-speed cache corresponding to the at least one data segment to be migrated to a target on the another level,
- reading remaining data, not previously copied from the lower-speed cache, from a source on the one level, and writing the remaining data to the target, and
- during the reading and the writing of the remaining data, allowing destages of updates corresponding to the at least one data segment to proceed to the lower-speed cache in lieu of abstaining from destaging the updates to the source.

2. The method of claim 1, further including changing a logical address of the at least one data segment to point to the target.

3. The method of claim 1, further including, subsequent to the reading and the writing of the remaining data, destaging the updates from either the higher and lower speed caches to the target.

4. The method of claim 1, further including, pursuant to copying the data resident in the lower-speed cache, identifying the resident data as having been copied.

5. The method of claim 1, further including, if at least one of the lower-speed cache is inaccessible and the lower-speed cache is filled with additional modified data to be subsequently destaged:

- initiating an alternate data migration mechanism whereby all data is read from the source and destaged to the target, while all destages of updates corresponding to the at least one data segment are withheld until the all data is destaged to the target.

11

6. The method of claim 1, further including performing initiating the data movement mechanism based on an observation of a heat metric.

* * * * *

12